
IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: Oehler et al.

Attorney Docket No.: NWISP001

Application No.: 09/932,456

Examiner: LEE, PHILIP C.

Filed: August 16, 2001

Group: 2152

Title: COMPUTER SYSTEM PARTITIONING
USING DATA TRANSFER ROUTING
MECHANISM

Confirmation No.: 3395

CERTIFICATE OF EFS-WEB TRANSMISSION

I hereby certify that this correspondence is being transmitted electronically through EFS-WEB to the Commissioner for Patents, P.O. Box 1450 Alexandria, VA 22313-1450 on July 10, 2007.

Signed: /Mia Mitchell-Haynes/
Mia Mitchell-Haynes

PRE-APPEAL BRIEF REQUEST FOR REVIEW

Mail Stop AF
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

Applicant requests review of the final rejection in the above-identified application. No amendments are being filed with this request.

This request is being filed with a Notice of Appeal.

The review is requested for the reasons stated below.

REMARKS

The Examiner has failed to establish a prima facie case of obviousness

In the final office action dated May 30, 2007, the Examiner rejected claims 1-4, 6-10, 12, 13, 16, 21, 23-26, 28, 31, 36-37, 39, and 40 under 35 U.S.C. 103(a) as being unpatentable over what the Examiner has referred to as Applicant's admitted prior art (AAPA) in view of U.S. Patent No. 5,303,383 (Neches). The Examiner also rejected claims 1, 15, 23, 30, 36, and 39 over AAPA and Neches in view of U.S. Patent No. 6,961,761 (Masuyama). The Examiner also rejected claims 1, 5, 11, 19, 22, 23, 27, 34, 36, 38, 39 and 41 over AAPA and Neches in view of U.S. Patent Publication No. US 2001/0037435 A1 (Van Doren). The Examiner also rejected claims 1, 14, 17, 23, 29, 32, 36, and 39 over AAPA and Neches in view of U.S. Patent No.

6,188,759 (Lorenzen). The rejections are respectfully traversed.

The primary issue for which this request has been filed relates to the fact that the Examiner has yet to provide any prior art references which teach or suggest a computer system which includes “a point-to-point transmission infrastructure” and “at least one partitioning processor for configuring the plurality of resources into a plurality of partitions, each partition comprising a subset of the plurality of resources and a portion of the point-to-point transmission infrastructure, *the portion of the point-to-point transmission infrastructure in each partition being distinct from and non-overlapping with the portion of the point-to-point transmission infrastructure in each other partition.*”

The Examiner has correctly pointed out that the Background of the present application describes a computer system having a point-to-point transmission infrastructure which interconnects a plurality of processors. The Examiner has also correctly pointed out that the Neches reference describes a partitioning scheme in which a network of processing modules interconnected by a switch fabric may be partitioned into “logically independent subsets” in which “[c]ommunication within one subset is prevented from interfering with communication in any other subset.” See column 35, lines 41-44. However, the Examiner’s mischaracterized Neches’ switch fabric as teaching the missing limitation referred to above when, in fact, it does not. In addition, the Examiner’s combination of Neches with the system described in the Background of the present application fails because of the incompatibility of Neches’ switch fabric with the point-to-point architecture described (i.e., dedicated physical links between processors).

Neches describes a multistage interconnect network (MIN) for interconnecting processors modules (PM). See Abstract and FIG. 2. It is clear from the figures and the description that network 14 is a switch fabric which provides multiple, redundant, and dynamically configurable paths between each PM 12 and any other PM 12. See, for example, column 7, lines 10-27. This is clearly distinguishable from the “point-to-point transmission infrastructure” which includes “dedicated physical links” between processors as described and claimed in the present application.

The references in Neches to “point-to-point” communication are clearly references to higher level communications among processors and refer to logical connections between processor modules rather than dedicated physical links as described and claimed in the present application. That is, as used in Neches, the term “point-to-point” refers to a communication protocol rather than the transmission infrastructure which, instead of being a point-to-point transmission infrastructure, is a hierarchical switch fabric. See, for example, column 6, lines 26-

51, and column 7.

Similarly, references in Neches to partitioning relate to the logical grouping of processor modules which result in groups of processor modules which communicate independently of each other via the same shared resource, i.e., the MIN. See, for example, column 35, line 40 et seq. That is, while communications between different “superclusters” of processor modules do not interfere, they must continue to use the same physical transmission infrastructure. Thus, it is impossible for the technique described in Neches to result in a partitioning schema in which “the portion of the point-to-point transmission infrastructure in each partition [is] distinct from and non-overlapping with the portion of the point-to-point transmission infrastructure in each other partition.” To the contrary, because Neches’ switch fabric is shared among the PMs 12, the transmission infrastructure employed by each of Neches’ superclusters *must* be overlapping with that used by other superclusters, i.e., it is the same physical infrastructure.

In addition, the Examiner’s combination of Neches with the description from the Background of the Invention fails to result in an operable system. That is, the multistage interconnect network (MIN) of Neches is not compatible with the system described with reference to Fig. 2 of the present application. Neches’ MIN is operable to directly route packets from every one of the input ports to every one of the output ports (see FIGs. 1 and 2 and the corresponding description). It is clear that in order for Neches’ system to operate, the MIN must operate as an undivided, shared resource among the processor modules. Thus, there is no way to partition the MIN in the manner claimed in the present application and have it remain operable.

By contrast, the point-to-point communication links of the system described in the Background of the Invention of the present application are *dedicated* links between the processing nodes which, in some cases, makes it necessary for indirect transmissions between two processing nodes (i.e., via an intermediate processing node).

Because Neches’ MIN is an indivisible switch fabric which is capable of directly connecting any one of the connected processor modules to any other one of the processor modules, it is incompatible with a point-to-point architecture as described in the present application. Therefore, the partitioning technique described in Neches is similarly irrelevant to the claimed invention. In view of the foregoing, all of the rejections referring to Neches should be withdrawn.

None of the Masuyama, Van Doren, or Lorenzen references makes up for the deficiencies described above. Therefore, in view of the fact that none of the art of record teaches a partitioning technique in which “the portion of the point-to-point transmission infrastructure in each partition being *distinct from and non-overlapping* with the portion of the point-to-point

transmission infrastructure in each other partition,” it is respectfully submitted that the rejection of the claims over Wilson should be withdrawn.

I am the attorney or agent acting under 37 CFR 1.34

Respectfully submitted,
BEYER WEAVER LLP

/Joseph M. Villeneuve/
Joseph M. Villeneuve
Reg. No. 37,460

P.O. Box 70250
Oakland, CA 94612-0250
(510) 663-1100

APPENDIX OF PENDING CLAIMS

1. (Previously Presented) A computer system, comprising:
a plurality of resources including a plurality of processors;
a distributed point-to-point transmission infrastructure for interconnecting the plurality of processors; and
at least one partitioning processor for configuring the plurality of resources into a plurality of partitions, each partition comprising a subset of the plurality of resources and a portion of the point-to-point transmission infrastructure, the portion of the point-to-point transmission infrastructure in each partition being distinct from and non-overlapping with the portion of the point-to-point transmission infrastructure in each other partition, the at least one partitioning processor being operable to configure the resources by writing to at least one of a plurality of routing tables associated with the processors according to a previously specified partitioning schema, each routing table representing dedicated physical links between an associated processor and other ones of the plurality of processors, the links corresponding to portions of the point-to-point transmission infrastructure.
2. (Original) The computer system of claim 1 wherein the plurality of resources further includes at least one of a memory device, a memory range, an I/O bus, I/O devices coupled to an I/O bus, and an interrupt mechanism for routing interrupts.
3. (Original) The computer system of claim 1 wherein the plurality of resources includes an I/O switch, the I/O switch having one of the routing tables associated therewith representing links between the I/O switch, at least one of the processors, and at least one I/O resource.
4. (Original) The computer system of claim 3 wherein the at least one I/O resource comprises at least one of an Ethernet device and a SCSI device.
5. (Original) The computer system of claim 1 wherein each routing table comprises a table of entries, each of selected ones of the entries associating an address of one of the resources with one of the processors and a link for connecting with the one of the processors.
6. (Original) The computer system of claim 1 wherein the distributed point-to-point transmission infrastructure comprises a coherent HyperTransport (cHT) infrastructure.
7. (Original) The computer system of claim 1 wherein the distributed point-to-point transmission infrastructure interconnects the processors using a ring topology.
8. (Original) The computer system of claim 1 wherein the distributed point-to-point transmission infrastructure interconnects the processors using a mesh topology.

9. (Original) The computer system of claim 1 wherein the distributed point-to-point transmission infrastructure directly connects each of the processors with every other one of the processors.

10. (Original) The computer system of claim 1 wherein the at least one partitioning processor comprises at least one of the plurality of processors interconnected by the distributed point-to-point transmission infrastructure.

11. (Original) The computer system of claim 1 wherein the at least one partitioning processor is separate from the plurality of processors interconnected by the distributed point-to-point transmission infrastructure.

12. (Original) The computer system of claim 11 further comprising a boot memory for facilitating initialization of the computer system, the boot memory having computer program instructions stored therein for facilitating operation of at least one of the plurality of processors as the at least one partitioning processor.

13. (Original) The computer system of claim 1 wherein the previously specified partitioning schema is generated in response to an event occurring during operation of the computer system.

14. (Original) The computer system of claim 13 wherein the event comprises one of initialization of the computer system, a failure of at least one of the resources, a change in operating load associated with at least one of the resources, passage of a period of time, use of particular software, and a change in available power resources.

15. (Original) The computer system of claim 1 further comprising at least one partitioning processor link for connecting the at least one partitioning processor with a user interface, and wherein the previously specified partitioning schema is specified by a user of the computer system via the user interface and the at least one partitioning processor link.

16. (Original) The computer system of claim 1 wherein the at least one partitioning processor is operable to generate the routing tables upon initialization of the computer system.

17. (Original) The computer system of claim 1 wherein the at least one partitioning processor is operable to alter the at least one of the routing tables during operation of the computer system.

18. (Canceled)

19. (Original) The computer system of claim 18 wherein at least one of the plurality of partitions comprising a functional subset of the plurality of resources.

20. (Canceled)

21. (Original) The computer system of claim 1 wherein the at least one partitioning processor comprises one partitioning processor.

22. (Original) The computer system of claim 1 wherein the at least one partitioning processor comprises more than one partitioning processor.

23. (Previously Presented) A computer implemented method for use in a computer system having a plurality of resources including a plurality of processors and a distributed point-to-point transmission infrastructure for interconnecting the plurality of processors, the method comprising configuring the plurality of resources into a plurality of partitions, each partition comprising a subset of the plurality of resources and a portion of the point-to-point transmission infrastructure, the portion of the point-to-point transmission infrastructure in each partition being distinct from and non-overlapping with the portion of the point-to-point transmission infrastructure in each other partition, the configuring of the resources being effected by writing to at least one of a plurality of routing tables associated with the processors according to a previously specified partitioning schema, each routing table representing dedicated physical links between an associated processor and other ones of the plurality of processors, the links corresponding to portions of the point-to-point transmission infrastructure.

24. (Original) The method of claim 23 wherein the plurality of resources includes an I/O switch, the I/O switch having one of the routing tables associated therewith representing links between the I/O switch, at least one of the processors, and at least one I/O resource.

25. (Original) The method of claim 24 wherein the distributed point-to-point transmission infrastructure comprises a non-coherent HyperTransport (ncHT) infrastructure.

26. (Original) The method of claim 23 wherein configuring the plurality of resources is achieved using at least one partitioning processor which comprises at least one of the plurality of processors interconnected by the distributed point-to-point transmission infrastructure.

27. (Original) The method of claim 23 wherein configuring the plurality of resources is achieved using at least one partitioning processor which is separate from the plurality of processors interconnected by the distributed point-to-point transmission infrastructure.

28. (Original) The method of claim 23 further comprising generating the previously specified partitioning schema in response to an event occurring during operation of the computer system.

29. (Original) The method of claim 28 wherein the event comprises one of initialization of the computer system, a failure of at least one of the resources, a change in operating load associated with at least one of the resources, passage of a period of time, use of particular software, and a change in available power resources.

30. (Original) The method of claim 23 further comprising receiving the previously specified partitioning schema as specified by a user of the computer system.

31. (Original) The method of claim 23 wherein writing to the at least one of the plurality of routing tables comprises generating the plurality of routing tables upon initialization of the computer system.

32. (Original) The method of claim 23 wherein writing to the at least one of the plurality of routing tables comprises altering the at least one of the routing tables during operation of the computer system.

33. (Canceled)

34. (Previously presented) The method of claim 33 wherein at least one of the plurality of partitions comprising a functional subset of the plurality of resources.

35. (Canceled)

36. (Previously Presented) A computer system, comprising:
a plurality of resources including a plurality of processors;
a distributed point-to-point transmission infrastructure for interconnecting the plurality of processors; and

at least one partitioning processor for configuring the plurality of resources into a plurality of partitions, each partition comprising a subset of the plurality of resources and a portion of the point-to-point transmission infrastructure, the portion of the point-to-point transmission infrastructure in each partition being distinct from and non-overlapping with the portion of the point-to-point transmission infrastructure in each other partition, the at least one partitioning processor being operable to configure the resources by enabling operation of at least one dedicated physical link between at least one of the plurality of processors and at least one other one of the plurality of processors according to a previously specified partitioning schema, the at least one link corresponding to a portion of the point-to-point transmission infrastructure.

37. (Previously Presented) The computer system of claim 36 wherein enabling operation of the at least one link comprises writing to at least one of a plurality of routing tables associated with the processors according to the previously specified partitioning schema.

38. (Previously Presented) The computer system of claim 36 wherein enabling operation of the at least one link comprises closing at least one switch associated with the at least one link according to the previously specified partitioning schema.

39. (Previously Presented) A computer implemented method for use in a computer system having a plurality of resources including a plurality of processors and a distributed point-

to-point transmission infrastructure for interconnecting the plurality of processors, the method comprising configuring the plurality of resources into a plurality of partitions, each partition comprising a subset of the plurality of resources and a portion of the point-to-point transmission infrastructure, the portion of the point-to-point transmission infrastructure in each partition being distinct from and non-overlapping with the portion of the point-to-point transmission infrastructure in each other partition, the configuring of the resources being effected by enabling operation of at least one dedicated physical link between at least one of the plurality of processors and at least one other one of the plurality of processors according to a previously specified partitioning schema, the at least one link corresponding to a portion of the point-to-point transmission infrastructure.

40. (Previously Presented) The method of claim 39 wherein enabling operation of the at least one link comprises writing to at least one of a plurality of routing tables associated with the processors according to the previously specified partitioning schema.

41. (Previously Presented) The method of claim 39 wherein enabling operation of the at least one link comprises closing at least one switch associated with the at least one link according to the previously specified partitioning schema.